Using a Hidden Markov Model for Resident Identification

Aaron S. Crandall
School of Electrical and Computer Science
Washington State University
Pullman, Washington 99164–2752, USA
Email: acrandal@wsu.edu

Diane J. Cook
School of Electrical and Computer Science
Washington State University
Pullman, Washington 99164–2752, USA
Email: cook@eecs.wsu.edu

***Abstract*—In smart home environments, it is highly desirable to know who is performing what actions. This knowledge allows the system to accurately build individuals' histories and to take personalized action based on the current resident. Without a good handle on identity, multi-resident smart homes are less effective when used for medical and assistive applications.**

Most smart home systems either have a single occupancy requirement, or rely on a wireless or video device to identify individuals. These requirements are too burdensome in some situations, which can limit the deployment of smart home technologies in environments that would derive benefits from them. This research work introduces the use of passive sensors and a Hidden Markov Model as a means to identify individuals. The result is a passive, low profile means to attribute individual events to unique residents.

For this work, two different pairs of individuals living in a smart home testbed are used to evaluate the tools. The data used is from unscripted, full time occupancy and annotated by the residents themselves for accuracy. Lastly, the Hidden Markov Model approach is compared and contrasted against a prior Naive Bayes solution on the same data sets.

I. Introduction

As smart home systems move towards commercialization and real world usage, the need to handle multi-resident situations becomes ever more important. To deal with multiple people, the smart home needs to address the problems of tracking and identifying the residents in some manner. Significant headway has been made in tracking individuals through spaces using wireless devices [1][2][3]. These devices can provide a rich source of often highly accurate information, but require that individuals maintain them. By moving to a passive and remote monitoring system, these maintenance issues can be avoided. In this work, the system relies on the Center of Advanced Studies in Adaptive Systems (CASAS) testbeds where only passive, non-intrusive sensors are deployed.

When it comes to activity recognition, there has been significant work using video data [4][5][6], motion sensor data [7][8], or other sources of information [9][10]. Some works have been published on dealing with an individual doing interleaved activities[11] and transferring learned activity models between spaces[12]. Many of the current technologies assume a single individual, or treat the group as a single entity[13], and need to be extended to handle the multiple resident problem in one way or another.

Given that the goal of this research project is to model and automate resident activity in multiple-resident intelligent environments, addressing the problem of how to identify each individual of the space is a root issue. There are numerous strategies that would ease the complexity of this task. For example, we could ask residents to wear portable devices that enable wireless tracking through the space [14][3]. In many situations, this is an impractical solution when individuals do not want to wear the device, forget to wear the device, or there is significant interference from other wireless sources. Similarly, capturing resident behavior with video cameras aids in understanding resident behavior even in group settings [15]. However, surveys with target populations have revealed that many individuals are adverse to embedding cameras in their personal environments [5]. As a result, our aim is to identify the individuals and their activities in an intelligent environment using only passive sensors.

Our first step is to design an algorithm that maps sensor events to the resident that is responsible for triggering a sensor event. This information will allow other tools to build their histories and profiles of resident behaviors. These knowledge bases can then be used to identify the individuals currently in the environment, predict their desires and interact with them on an individual basis. Some previous works have focused on passive multi-resident systems [16], and give some indication of techniques that have succeeded on real-world data sets for activity recognition [17].

The solutions used in this work revolve around using very simple passive sensors, such as motion, contact, door sensors, appliance interaction and light switches to give a picture of what is transpiring in the space. These information sources offer the benefits of being fixed, unobtrusive and robust devices.

Since smart home research has the ultimate goal of being deployable in real-world environments, seeking solutions that are as robust as possible is always a factor in the systems we engineered. With that in mind, building an entirely passive solution gives the advantage of keeping the technology physically separate from the inhabitants while they go about performing their daily routines. This lets the smart home feel as ”normal” as possible to the residents and their guests. To date, every resident in a CASAS testbed has noted that they rapidly forget that the system is installed and running within a few weeks of occupying the space. By reducing the profile of the new devices as much as possible, people should be less effected by the technology that surrounds them.

After working with the data from the CASAS testbeds, our researchers and annotators noted several key things. First, was that given experience, it was possible to quickly differentiate between residents through viewing the data alone. Second, the identification of the individual was done using only a handful of sensors and some contextual events. Lastly, this skill was not limited to only a handful of residents and was common among our group of annotators. Taken together, it was hypothesized that a classification tool given a corpus of annotated data and a series of events to provide context would do a good job at identifying the current individuals in the smart home space. Of the available algorithms, the Hidden Markov Model (HMM) has the features that would test this hypothesis well.

To test this hypothesis, a HMM was implemented and tested on two real world data sets from one of the CASAS testbeds. The HMM is then compared and contrasted against a Naïve Bayesian Classifier (NBC), as was published in[18]. These two models have significant differences in their behavior, and both have been proven as robust on similar kinds of problems. This allows for different perspectives and insights into the same problem.

The HMM solution proposed in this work offers the advantage of using previous behavioral data collected from the set of known residents, but without requiring significant additional day to day actions to be performed by the residents. This historical behavior is used to train the learning algorithm for use in future real-time classification of the individuals and can be updated over time as new data arrives.

Both algorithms used in this work are robust and scalable, as well as real time capable. The end goal is to be able to run such systems on limited hardware in a home, so both memory and execution efficiency were considerations when choosing these tools. We validate our hypothesis using data collected in a real smart home environment with volunteer participants living in the space full time.

II. Data Gathering Environment

The testbed used for both data sets presented in this paper were gathered at one of Washington State University's CASAS testbeds. This particular installation is a two level apartment, shown with sensor placement in Figure 1. The apartment is inhabited by volunteer students who live there full time. None of this data was scripted or guided, and it was annotated for accuracy by the residents themselves.

The CASAS testbeds are designed to capture both spatial and temporal events through a variety of sensors installed throughout the space. The most common sensor is a passive infrared motion detector placed on the ceilings and pointed straight down. These sensors have their view occluded to only "see" a roughly 1.5m x 1.5m section of the floor below them. By arranging the sensors in a roughly uniform pattern

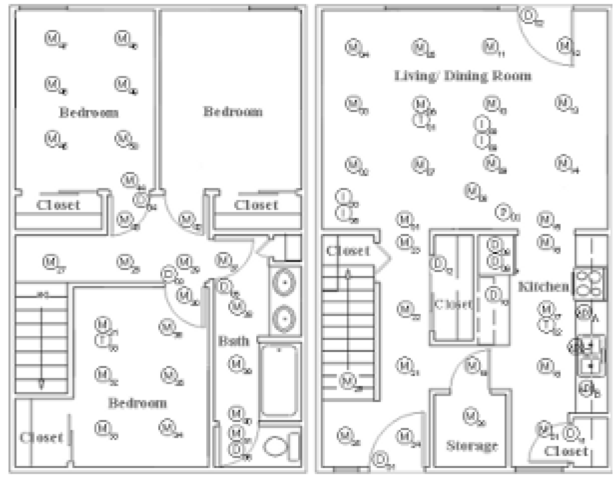

Fig. 1. Floor plan for WSU CASAS student apartment testbed, named Kyoto.

throughout the space the system gathers information about where and when activity is taking place. The view of the sensors is not perfect, and they sometimes overlap anywhere from 5cm to 30cm. These are commodity home security sensors with an added daughter board to make them easily accessible in large numbers to a single computer. They are capable of detecting movements as small as reaching for a mouse when working at a computer.

Additionally, the space has a number of other types of sensors throughout. There are open/close sensors on doors and cabinets, water flow sensors on the kitchen sink, power line control light switches and power usage to monitor the stove. The testbed has more kinds of sensors installed for other studies, but that data was left out of this work. These other types of sensors capture details about behavior, especially in the kitchen.

By being able to capture spatial and temporal information throughout the home, the algorithms are able to detect the unique behaviors of the various residents. All of these sensors require no direct interaction with the residents. This low profile and passive architecture allows the residents to live as they normally do, while giving the computer the opportunity to determine the current activities and identities of those in the space.

III. Data Representation

The data used for this study is available from the CASAS shared data sets repository[19] and comes in the form of a quintuple, as shown in Table I.

The first four fields are determined by the data collection infrastructure. The last field is added by manual annotation by the residents and researchers. This annotated class is done on an event by event basis, where each individual event is tagged with which known individual caused that event.

This particular testbed has been in operation for over two years and has had four different sets of residents to date. In this study two different pairs of residents were chosen for the data sources. The earlier pair, known as "B&B" helped to annotate

TABLE I
DATA SET QUINTUPLE DESCRIPTION

Field	Example
Date	2007-05-19
Time	10:13:19.3423
Serial	M001
Message	ON
Annotated Class	R1 (unique resident identifier)

TABLE II
DATA SET EXAMPLE

Date	Time	Serial	Message	ID
2009-02-02	22:29:18.31887	M35	ON	R1
2009-02-02	22:29:18.64348	M45	OFF	R2
2009-02-02	22:29:22.615809	M47	ON	R2
2009-02-02	22:29:22.913059	M46	ON	R2
2009-02-02	22:29:28.184089	M35	OFF	R1
2009-02-02	22:29:28.298049	M47	OFF	R2

three days worth of data. They were both present all or part of each day in the testbed.

The second data set, known as ”TwoR”, covers a much longer two months of residency. Again, both residents helped to annotate the data and they occupied the space most days during their tenure. There was no scripted or guided behaviors or studies and they used all facilities of the home.

The B&B data set consists of 10,193 events and the TwoR set is 136,504 events. For an example of the data used by the classifiers, see Table II. The goal for the classifiers is to take a subset of the given data as training, then accurately classify the remaining data set on an event by event basis as to who caused which event. A good classifier will exploit both the temporal and physical information available from the training set, though no explicit information regarding the layout of the sensors was provided to these tools. For the Naïve Bayes classifier, this is done through a likelihood of any individual event being caused by a given resident, while the Hidden Markov Model can take into account the whole series of events to classify them.

IV. THE CLASSIFIERS

The two classifiers used in this research were a Hidden Markov Model and a Naïve Bayesian Classifier. The NBC has been applied to a number of smart home problems[20] with significant success. It has been explored as a means to identify individuals in smart home spaces[18] and contrasted against a Markov Model solution[21] as well. The NBC is used as an established baseline to compare the capabilities of the HMM against.

The new algorithm introduced in this paper is a classic HMM and is defined as such and shown in Figure 2:

- The hidden states represent the possible residents in the data set.

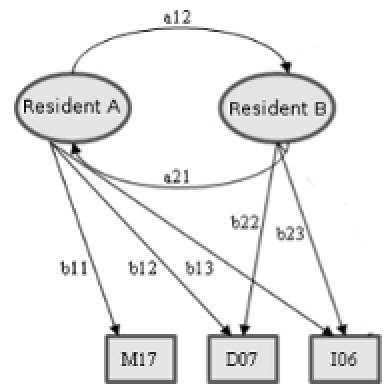

Fig. 2. HMM architecture of hidden states, transitions and observations.

TABLE III
OVERALL ACCURACIES FOR BOTH ALGORITHMS

	B&B	TwoR
NBC	93.3%	89.3%
	σ = 0.0014	σ = 0.0016
HMM	94.0%	90.2%
	σ = 0.0034	σ = 0.0021

- The possible observations are the sensors in the space combined with their message values. For example, a single motion detector can emit both an ”ON” and ”OFF” message, so the resulting HMM has observations for both states.
- The transition probabilities represent the likelihood the next event is from the same person or a different one.
- The emission probabilities are the likelihood that a given resident causes any given sensor+message combination.

The resulting classifiers were tested using 3-fold cross validation. The data sets were split into three groups by days. Each classifier was trained on two out of three days and tested on the remaining days. The results from all three run permutations were averaged together for an overall accuracy. Additional statistics showing the behavior of the classifiers and the data sets were gathered for insights into the capabilities of the tools.

V. RESULTS

Both algorithms performed well on both data sets. The overall accuracy values are shown in Table III. The HMM performed slightly better, though not significantly so.

Given the complexity of the data with multiple residents and no given structure to their behavior, the highly accurate results from both algorithms attests to their robustness. The NBC accuracy dropped significantly from earlier works[18] based on the same technological platform. This was likely due to the greater overlap in behavior compared to the simpler laboratory setting used previously.

Overall, the HMM results are promising. The initial hypothesis that drawing on additional contextual information across events would allow an algorithm to better differentiate between

TABLE IV
EXAMPLE HMM BEHAVIOR PATTERN

Event Number	Annotated Class	Chosen Class	Success
1	R1	R1	SUCCESS
2	R1	R1	SUCCESS
3	R2	R1	FAIL
4	R2	R2	SUCCESS
5	R2	R2	SUCCESS

individuals seems to be supported by the overall accuracy results.

When analyzing the actual pattern of classification, the behavior of the HMM is more complex than the NBC. As the events arrive, it takes the HMM zero or more additional events to determine to whom the new events belong to. For an example of this behavior, Table IV shows a small snippet of events as classified by the HMM. The left column is the event number, the second column is annotated resident value for the event, the third is what the algorithm determined, and the last column is the success or fail on the given event. This snippet has a transition from R1 to R2 at event #3. The HMM takes until event #4 before it has enough evidence to change states. This situation, where the events change from one resident to another, has been termed a ”transition” and is an important feature in how the HMM behaves.

By the overall accuracy metric, this is a score of 4/5, or 80% accuracy. What is most interesting about this series is that the events arriving at the computer are initially from R1, then change to R2 at some point. With NBC, it takes every event with its message in isolation, so there is no previous context to consider so it is either correct or incorrect with every event. Conversely with the the HMM, there is now a possibility of a transition time as the evidence that the new events are from a different person accumulates. The concern is that this transition time would significantly impact the effectiveness of the HMM as a tool for identification.

To determine how much this transition error is effecting the HMM, several statistics were gathered from the final classifications. The first was the total number of times in the event stream the annotated resident value switches from one to another. This is an indication of the data complexity. If the number of transitions increases it indicates more simultaneous occupancy of the space, which can be more difficult for the HMM to accurately classify.

The hypothesized relationship between the rate of transitions in the data set and the final accuracy was not borne out by the results, as shown in Figure 3. The transition rate line was expected to trend upwards, opposite the overall accuracies of the classifier. Instead it is found to trend downward between the two data sets presented here. On further inspection, it is not merely the number of transitions that effects the overall accuracy, but more commonly the physical locality of the residents during those transitions. If the people are ”close” together in the space, there is less evidence in the emission probabilities that the HMM should change its hidden state.

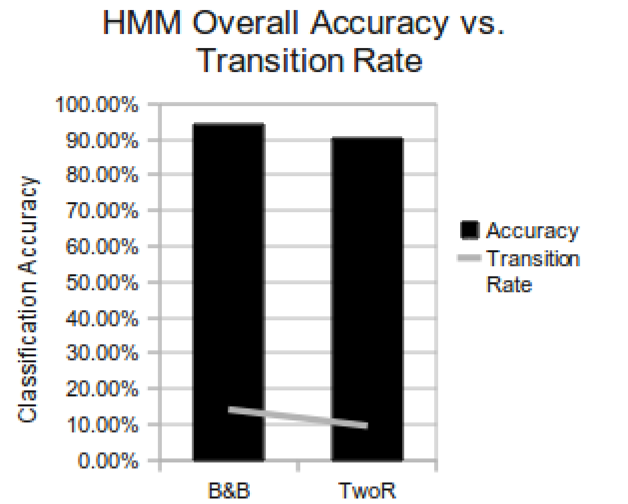

Fig. 3. HMM overall accuracy for each data set, with the data sets' comparative transition rate. The transition rate was expected to trend opposite to the overall accuracy instead of with it.

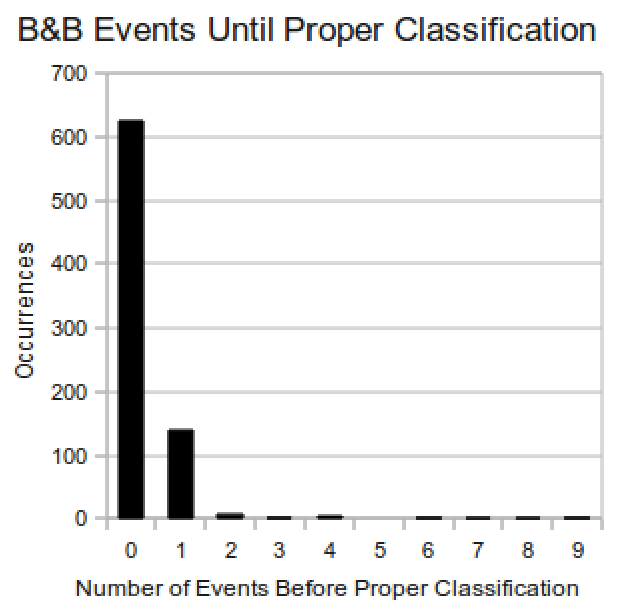

Fig. 4. HMM lag lengths in proper classification after a transition between residents in the B&B data set.

Essentially, the physical layout of the sensor space is encoded to some degree within the HMM itself. If there is no separation between individuals in the sensor space, then there is less evidence that events are coming from a different person and the HMM lags before changing state.

As a measure of how much this lag in transition impacts the behavior of the algorithm, some numbers about the delay were gathered. The primary number to look at is the average number of events after a transition before the HMM properly changes to accurately classify the resident. To find this value, the results were processed for the length of the lag in the transition on each data set. Figures 4 and 5 show the total occurrences of lag lengths (zero or more), grouped by length until proper classification for each data set.

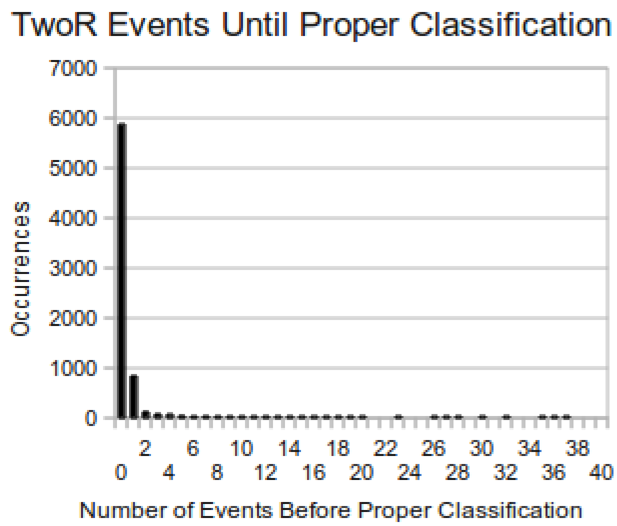

Fig. 5. HMM lag lengths in proper classification after a transition between residents in the TwoR data set.

TABLE V
HMM AVERAGE LAG LENGTH FOR BOTH DATA SETS. AN AVERAGE OF ZERO IS PERFECT TRANSITION CLASSIFICATION.

Data Set	Average Lag	Standard Deviation (σ)
B&B	0.19	0.80
TwoR	0.38	2.17

The first column in Figures4 and 5 represents the count of transitions in each data set where the HMM changes state properly on the very first event after a transition. In these cases the HMM properly transitions from one resident to another with only the very first event as evidence. The rest of the columns are instances where there is one or more events improperly classified before the HMM transitions properly. This ”lag” in the HMM after transitions in the data are a significant portion of the HMM's overall error.

Table V shows the average length of the lag in the HMM transition for each data set. An average of zero would mean that it has no lag whatsoever on the given data set, leading to perfect classification during transitions. The lower average lag for the B&B data set is in line with the overall higher accuracy. This indicates that the HMM was able to use the evidence to more rapidly transition between residents based upon their behavior in the sensor space. The TwoR residents were admittedly more social than the B&B residents, and spent more time together during their stay in the testbed. By spending more time in close proximity, the resolution of the sensor network had more trouble providing evidence for the HMM to determine who was whom during the close interactions and the overall accuracy suffered.

The other sign the the TwoR residents were more often interacting during the time of this data gathering is the longer lengths of the HMM's transition lag. With the B&B data set, there were very few instances where the HMM was not able to properly transition within one or two events. This indicates that the residents were most often physically separated in the testbed space. The very long lag occurrences induced by the TwoR were observed to be when the two residents were doing activities like cooking or homework together. In those cases, the HMM was unable to differentiate between the residents for quite some time.

TABLE VI
HMM NON-TRANSITION CONFUSION ERROR RATES.

Data Set	Confusion Error
B&B	3.2%
TwoR	6.1%

Another source for the error in classification is when the HMM chooses the incorrect class, but there was no actual transition to another resident. In this case the algorithm is truly confused, and this error type is more akin to the type of error in the NBC. The total error rate for this kind of mis-identification are summed up in Table VI. The higher rate for the TwoR data set indicates that these two individuals had more behavior that was similar to one another than the two people in the B&B data, which again contributes to the lower overall accuracy on the TwoR set.

VI. CONCLUSION

The HMM has managed to reduce the kind of error that the NBC was generating from about 6.5% to 3.2% on the B&B data and from 10.7% to 6.1% on the TwoR data, but introduced an additional lag in identification as new evidence arrives when the residents are in relative proximity to each other. The reduction in general confusion indicates that our original hypothesis about additional context being valuable to identification holds true. The HMM is able to take into account a series of events to do a better job at identifying a given individual based on their behavior alone. Continued work into which sensor platforms and data features should give the HMM more evidence to differentiate between residents. This will reduce both the general confusion and to shorten the average number of events the algorithm needs to transition between individuals in multi-resident situations.

As with earlier works, it has been demonstrated that incorporating additional features about the data, such as temporal length or time of occurrence for events can improve these kinds of accuracies. This HMM algorithm did not have the advantage of either one of these kinds of temporal features and could be improved with additional modifications to the data features.

The HMM tool also provides more insights into the reasons for error and the behavior of the residents. With the NBC it was difficult to tell why it made a given choice at run time. By analyzing the series of classifications by the HMM next to the sensor map, the researchers could determine what behaviors were easy or difficult for the HMM to classify. It was then easier to algorithmically detect the reasons for success or failure within the final results. The opaque nature of the NBC made similar analysis much more troublesome.

While this study did not have data sets with large numbers of residents, the classifiers are likely to scale differently. The HMM should scale against the number of residents better than the NBC by using the greater context available in the model to make fine differentiations between individuals. It also requires less tuning to make high quality classifications in these complex environments, which is a good sign for its ability to be deployed in real-world situations.

Overall, the researchers felt that the HMM was easier to understand and was less erratic in its behavior. This was corroborated by the lower general confusion in classification when there were no transitions between residents occurring. From experience and empirically, the HMM was significantly more stable then the NBC when a single resident was present and more understandable even when incorrect in complex situations.

VII. Future Work

Identifying, or at least classifying, individuals in smart home environments will continue to be an issue for some time. It is intrinsically linked to the definition of "behavior" for classification, the available information from the installed sensor system and the amount of annotated data available. Continued work into choosing the right sources, types and features from the sensor layer will need to guide which algorithms are used to identify individuals as they live in the space.

In the future, newer classifiers should be applied to the problem. The HMM is a generative model often used for classification, but tools such as conditional random fields might do a better overall job at dealing with this complex problem.

Some tools that might assist greatly include the introduction of passive biometrics to give stronger hints as to the current resident, or using a much more diverse set of sensors that are geared towards sensing personalized behavior and not just activities.

Acknowledgments

This work is supported by NSF grant IIS-0121297.

References

[1] P. Bahl and V. Padmanabhan, "Radar: An in-building rf-based user location and tracking system," in *Proc. IEEE Infocom*. IEEE CS Press, 2000, pp. 775–784.

[2] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location support system," *Proc. 6th Ann. Int'l Conference on Mobile Computing and Networking*, pp. 32–43, August 2000.

[3] J. Yin, Q. Yang, and D. Shen, "Activity recognition via user-trace segmentation," Jan 2007.

[4] L. Snidaro, C. Micheloni, and C. Chivedale, "Video security for ambient intelligence," *IEEE Transaction on Systems, Man and Cybernetics, part A*, vol. Vol. 35, pp. 133–144, 2005.

[5] S. S. Intille, "Designing a home of the future," *IEEE Pervasive Computing*, vol. vol. 1, pp. 80–86, Apr./June 2002.

[6] W. Feng, J. Walpole, W. Feng, and C. Pu, "Moving towards massively scalable video-based sensor networks," in *Proc. Workshop on New Visions for Large-Scale Networks: Research and Applications*, 2001.

[7] C. Wren and E. Tapia, "Toward scalable activity recognition for sensor networks," *Hazas, M., Krumm, J., Strang, T. (eds.) LoCA 2006*, vol. 3987, pp. 168–185, 2006.

[8] V. Jakkula and D. Cook, "Using temporal relations in smart environment data for activity prediction," in *Proceedings of the 24th International Conference on Machine Learning*, 2007.

[9] U. Naeem and J. Bigham, "Activity recognition in the home using a hierarchal framework with object usage data," *Journal of Ambient Intelligence and Smart Environments*, pp. 335–350, 2009.

[10] R. J. Orr and G. D. Abowd, "The smart floor: a mechanism for natural user identification and tracking," in *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*. New York, NY, USA: ACM, 2000, pp. 275–276.

[11] G. Singla and D. J. Cook, "Interleaved activity recognition for smart home residents," in *Proceedings of the 5th International Conference on Intelligent Environments*. Amsterdam, The Netherlands: IOS Press, 2009, pp. 145–152.

[12] P. Rashidi and D. J. Cook, "Transferring learned activities in smart environments," in *Proceedings of the 5th International Conference on Intelligent Environments*. Amsterdam, The Netherlands: IOS Press, 2009, pp. 185–192.

[13] B. El-Desouky and H. Hagras, "An adaptive type-2 fuzzy logic based agent for multi-occupant ambient intelligent environments," in *Proceedings of the 5th International Conference on Intelligent Environments*. Amsterdam, The Netherlands: IOS Press, 2009, pp. 257–266.

[14] J. Hightower and G. Borriello, "Location systems for ubiquitous computing," *Computer*, vol. 32, no. 8, pp. 57–66, August 2001.

[15] J. K. et al., "Multi-camera multi-person tracking for easy living," in *Proc. 3rd IEEE Int'l Workshop Visual Surveillance*. Piscataway, N.J.: IEEE Press, 2000, pp. 3–10.

[16] D. Cook and S. Das, "How smart are our environments? an updated look at the state of the art," *Journal of Pervasive and Mobile Computing*, 2007.

[17] C. Lu, Y. Ho, and L. Fu, "Creating robust activity maps using wireless sensor network in a smart home," in *Proceedings of the 3rd Annual Conference on Automation Science and Engineering*. Scottsdale, AZ: IEEE, September 2007.

[18] A. S. Crandall and D. J. Cook, "Attributing events to individuals in multi-inhabitant environments," in *Proceedings of the 4th International Conference on Intelligent Environments*. Amsterdam, The Netherlands: IOS Press, 2008.

[19] CASAS, "Smart home publicly shared datasets," World Wide Web, March 2010. [Online]. Available: http://ailab.wsu.edu/casas/datasets.html

[20] E. M. Tapia, S. S. Intille, and K. Larson, *Activity recognition in the home setting using simple and ubiquitous sensors*, ser. Proc. PERVASIVE. Berlin / Heidelberg: Springer-Verlag, 2004, vol. LNCS 3001, pp. 158–175.

[21] A. S. Crandall and D. J. Cook, "Resident and caregiver: Handling multiple people in a smart care facility," in *AI in Eldercare: New Solutions to Old Problems*. Menlo Park, California, USA: AAAI Press, 2008, pp. 39–47.